



Defence Research and
Development Canada

Recherche et développement
pour la défense Canada



Adding a Capability to Extract Sentiment from Text using HanDles

*Simon Dennis
Benjamin Stone
Jihun Hamm
Ohio State University*

*Peter Kwantes
DRDC Toronto*

Defence R&D Canada
Technical Memorandum
DRDC Toronto TM 2012-063
May 2012

Canada

Adding a Capability to Extract Sentiment from Text using HanDles

Simon Dennis
Benjamin Stone
Jihun Hamm
Ohio State University

Peter Kwantes
DRDC Toronto

Defence R&D Canada – Toronto

Technical Memorandum

DRDC Toronto TM 2012-0632012-063

May 2012

Principal Author

Simon Dennis

Contractor, Ohio State University

Approved by

Original signed by Keith Stewart

Keith Stewart

Head, Socio-Cognitive Systems Section

Approved for release by

Original signed by Dr. Joseph V. Baranski

Joseph Baranski

Chair, Knowledge and Information Management Committee; Acting Chief Scientist

[Include the sponsor of the work or a reference to a thrust or work unit, when deemed appropriate by author or CSA; relevant patent number(s), relating to protected intellectual property, should be noted. If there is no relevant information for this document, delete this text.]

In conducting the research described in this report, the investigators adhered to the policies and procedures set out in the Tri-Council Policy Statement: Ethical conduct for research involving humans, National Council on Ethics in Human Research, Ottawa, 1998 as issued jointly by the Canadian Institutes of Health Research, the Natural Sciences and Engineering Research Council of Canada and the Social Sciences and Humanities Research Council of Canada.

© Her Majesty the Queen in Right of Canada, as represented by the Minister of National Defence, 2012

© Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2012

Abstract

HanDles is a document visualization tool developed by Ohio State University for DRDC Toronto. One aspect of documents that might be of interest to analysts is the extent to which they express positive or negative opinion or sentiment toward some issue or group. In this report, we describe how HanDles was extended to include the ability to classify documents as containing predominantly positive or negative sentiment. To do so, we trained the semantic model underlying HanDles' understanding of the document collection to distinguish positive from negative documents. Our tests of the system suggested that its ability to discriminate positive from negative documents would be greatly improved by selecting a training collection that is similar in nature and content to the documents that will be evaluated in operational settings.

Résumé

HanDles est un outil de visualisation de documents conçu par l'Ohio State University pour RDDC Toronto. Une caractéristique des documents qui peut s'avérer intéressante pour les analystes est l'importance de l'opinion positive ou négative que dégagent ces documents à l'égard de certaines questions ou de certains groupes. Dans le présent rapport, nous décrivons comment nous avons amélioré HanDles afin qu'il prenne en charge la classification de documents selon la prédominance, dans leur contenu, de sentiments positifs ou négatifs. Pour ce faire, nous avons formé la compréhension du modèle sémantique sous-jacent à HanDles quant au recueil de documents utilisé afin qu'il soit en mesure de distinguer les documents positifs de ceux qui sont négatifs. Nos essais du système nous poussent à croire qu'il est possible d'accroître considérablement sa capacité à différencier les documents selon le sentiment qui s'en dégage en choisissant un recueil de formation dont la nature et le contenu ressemblent à ceux des documents qui seront évalués dans un contexte opérationnel.

This page intentionally left blank.

Executive summary

Adding a Capability to Extract Sentiment from Text using HanDles:

**Peter Kwantes; Benjamin Stone; Jihun Hamm; Peter Kwantes; DRDC Toronto
TM 2012-0632012-063; Defence R&D Canada – Toronto; May 2012.**

Introduction or background: HanDles is a document visualization tool developed for DRDC Toronto as part of ARP project 15ah. In this project, HanDles was augmented with the capability to classify documents as expressing either positive or negative sentiment. The capability was added to the tool so that it could be used in Influence Operations contexts in which analysts want to measure the extent to which issues or groups of interest are viewed favourably or unfavourably. As a test case, we trained HanDles to distinguish good and poor film reviews, and then tested it three times to see how well it classified documents. The first test was conducted on reviews of the Amazon Kindle. The second test was run on text segments of the original training set of movie reviews, and finally, it was tested on a set of movie reviews that it had not seen before.

Results: In general, HanDles did a poor job detecting the sentiment associated with the reviews of the Amazon Kindle. We attribute the poor performance to the fact that movie and product reviews discuss different issues, and as such, there will be limited similarity in the two classes of document. Not surprisingly, HanDles did a good job classifying text segments of the original training set. However, the finding demonstrated that, unlike many other sentiment analysis tools that only classify text at the whole-document level, HanDles can be used effectively to extract the issues being discussed within documents, and assign sentiment to those. For example, a review of a film might be classified as negative overall, but HanDles can determine that, for example, the *acting* was good, but the *directing* was poor. Finally, when we tested HanDles on a new set of movie reviews it had not seen before, it performed with 93.3% accuracy.

Significance: The results of our trial suggest that HanDles could represent a powerful tool for extracting sentiment and other kinds of higher-level properties from reports or intercepted media. What will be vital however, is that the system first be trained on the right kinds of documents. In other words, there must be some similarity between the documents used during training and those used in the operational context in order for it to work properly.

Future plans: To transition HanDles into operational use, DRDC and CF stakeholders must decide on what the most appropriate documents are for training the system to tell the difference between positive and negative opinion in text. Once a class of document has been decided, we can proceed to train it and trial the system in a more realistic context.

Sommaire

Adding a Capability to Extract Sentiment from Text using HanDles:

**Peter Kwantes; Benjamin Stone; Jihun Hamm; Peter Kwantes DRDC Toronto
TM 2012-0632012-063 ; R & D pour la défense Canada – Toronto; mai 2012.**

Introduction ou contexte : HanDles est un outil de visualisation de documents conçu pour RDDC Toronto dans le cadre du PRA 15ah. Dans ce projet, nous avons ajouté dans HanDles la capacité de classifier les documents selon l'impression positive ou négative qui s'en dégage. Nous avons apporté cette amélioration afin de pouvoir utiliser cet outil dans le cadre d'opérations d'influence au cours desquelles des analystes souhaitent mesurer l'importance des opinions favorable et défavorable envers des problèmes ou groupes d'intérêt donnés. Au cours du scénario d'essai, nous avons formé HanDles afin qu'il fasse la distinction entre bonnes et mauvaises critiques de film. Nous avons exécuté ce scénario à trois reprises dans le but de vérifier à quel point il parvient à classifier correctement les documents. Le premier essai portait sur des critiques de l'appareil Kindle d'Amazon. Le deuxième se concentrait quant à lui sur des extraits de textes tirés du recueil de critiques de film de formation. Finalement, nous avons effectué le troisième essai à l'aide d'un recueil de critiques de film qu'HanDles n'avait jamais traité auparavant.

Résultats : Dans l'ensemble, HanDles ne parvient pas vraiment à détecter le sentiment associé aux critiques du Kindle. Nous attribuons ce piètre rendement au fait que les critiques de films et d'appareils traitent des questions différentes, et donc qu'il n'y a que peu de similarités entre ces deux catégories de documents. Comme nous nous y attendions, il a effectué un bon travail quant à la classification des segments de texte provenant du recueil de formation. Cependant, nos résultats démontrent que, contrairement à de nombreux autres outils d'analyse d'impressions qui n'effectuent que la classification de textes complets, HanDles est en mesure d'extraire correctement les points traités dans les documents et de leur attribuer un sentiment. Par exemple, dans le cas d'une critique de film classifiée comme étant négative dans l'ensemble, il peut déterminer que le *jeu des acteurs* était bon, mais la *réalisation*, médiocre. Finalement, lorsque nous avons mis HanDles à l'essai à l'aide d'un recueil de critiques de films qu'il n'a jamais traité auparavant, il a classifié les documents avec une exactitude de 93,3 %.

Signification : Les résultats obtenus lors de nos essais nous portent à croire que l'outil HanDles pourrait constituer un puissant outil permettant d'extraire les sentiments émanant de rapports ou de médias interceptés. Il est cependant primordial de former le système à l'aide des bons types de documents. En d'autres mots, pour que cet outil fonctionne adéquatement, il doit y avoir certaines similarités entre les textes employés au cours de la formation d'HanDles et ceux traités dans un contexte opérationnel.

Plans pour l'avenir : Pour effectuer la transition d'HanDles vers un environnement opérationnel, les intervenants de RDDC et des FC doivent décider quels sont les documents les plus appropriés à la formation du système afin qu'il soit en mesure de bien discerner les opinions positives de celles négatives. Lorsqu'ils auront choisi une catégorie de texte, nous pourrions procéder à la formation et à la mise à l'essai du système au sein d'un environnement plus réaliste que ceux utilisés jusqu'à présent.

Table of contents

Abstract	i
Résumé	i
Executive summary	iii
Sommaire	iv
Table of contents	v
List of figures	vi
Acknowledgements	vii
1 Introduction.....	1
1.1 Purpose of the work.....	1
1.1.1 Integrating OMSA into the HanDles Tool.....	1
1.1.2 The Classifier.....	2
1.1.3 Evaluation	3
2 Conclusion and Recommendations.....	12
References	13
List of symbols/abbreviations/acronyms/initialisms	15

List of figures

Figure 1 Documents View, 67 “Star Wars” movie reviews have been split into document size 50, resulting in 1214 handles documents.	3
Figure 2 The Space View has been organized by Negative and Positive Sentiment.....	4
Figure 3 Negative Sentiment and “George Lucas”.	5
Figure 4 Negative Sentiment and “Obi Wan Kenobi”.....	6
Figure 5 Positive Sentiment and “Special Effects”.	7
Figure 6 The potential problem of context for Sentiment Classification	7
Figure 7 The problem of “non Star Wars” reviews that mention “Star Wars”.....	8
Figure 8 HanDles map of 30 film reviews from IMDB users	10
Figure 9 Alignment of the film titles when positive and negative handles are pulled apart and 'pumped'	10
Figure 10 HanDles misclassified two reviews: Gran Torino and Jack and Jill.	11

Acknowledgements

This work was conducted as a Work Breakdown Element (15ah02) of Applied Research Project 15ah.

This page intentionally left blank.

1 Introduction

1.1 Purpose of the work

Traditional information retrieval mechanisms focus on the content of a document--that is, the issues that the document raises. Just as critical in many intelligence and Influence Operations contexts is the stance that the author takes with respect to those issues, or how that author feels about them. The area of automated opinion mining and sentiment analysis (OMSA) uses natural language processing and machine learning techniques to classify documents into positive and negative classes. The objective is to facilitate an analyst's task in identifying critical documents to study from large collections and to provide a global view of the sentiment expressed across the documents.

The purpose of this work was to introduce a sentiment analysis mechanism into the HanDles document visualization tool.

1.1.1 Integrating OMSA into the HanDles Tool

HanDles provides a search interface to a document set and three visualizations to allow participants to quickly assimilate large document collections. The first of these visualizations is a typical results set as would be returned from an Internet search engine such as Google or Bing. In addition, however, HanDles provides a set of automatically generated tags known as *handles*, that allow the user to quickly select subsets of documents of interest. The second visualization plots a projection of a semantic space derived from the documents. Both handles and documents are plotted with proximity coding for semantic similarity. Rather than provide a static display, however, handles and documents can be dragged to interactively modify the view of the space afforded to the user. Finally, a timeline view plots the popularity of different handles as a function of time.

To integrate OMSA into the tool two special handles titled "Positive Sentiment" and "Negative Sentiment" were added. The user can choose to have these handles displayed by clicking on a sentiment link in the Search Options box on the main results page of the tool. The positive and negative handles then become available in all three views of the interface. In the document view, they allow the user to quickly highlight either the positive or negative documents, so that they can scan those quickly. This view also provides counts of the positive and negative documents, allowing the user to rapidly assess the general sentiment towards the topic for which they searched in the document set.

In the space view, handles expressing content issues can be arranged around the screen. For instance, if you were a market researcher examining reviews of the Amazon Kindle you might arrange issues like page turning, battery life, customer support, and other aspects around the screen. Selecting one of the sentiment-labeled handles then highlights the documents that are classified with that sentiment. The distribution of those highlighted documents makes it visually apparent how sentiment is distributed across those issues. In the Kindle example, this would allow you to quickly determine not only that customer support is an issue of interest to the reviewers, but also whether they were positively or negatively disposed towards Amazon's

support. The way in which handles cluster also highlights the hot spots that might require more detailed study. At any time, a user can click directly on the dot representing a document to deepen their investigation.

In the timeline view, the positive and negative handles can be used to trace how sentiment has changed as a function of time. This function is critical in operations in which one is attempting to change the sentiment expressed by authors to monitor the success of those programs.

In many cases, operating at the level of entire documents (news articles, movie reviews etc.) is too coarse. Within a document many issues may be raised and the sentiment expressed towards each issue may not be the same. A movie reviewer might have liked the plot, but think that the acting left something to be desired, for instance. To capture this, HanDles allows one to divide documents into subdocuments of fixed length. The main results screen provides a box in which one can specify how many words will be contained in these document fragments.

1.1.2 The Classifier

The ability to effectively use the sentiment mechanisms deployed in HanDles tool relies on the accuracy of the sentiment classifier. There has been a recent surge in work on sentiment classification, based in part on the provision of large sets of documents labeled with sentiment. In particular, Pang and Lee (2004) provided a movie review dataset consisting of 1000 positive and 1000 negative reviews crawled from the Internet Movie Database (IMDB) movie archive, with an average length of 30 sentences.

The algorithm we used to classify documents into positive and negative classes is called a Support Vector Machine (SVM). Simply put, an SVM examines a large collection of documents that have already been classified as positive or negative, and tries to work out the function that best differentiates them. After training, the SVM uses the function it has decided upon to classify new documents as being either positive or negative in tone or sentiment.

To try to achieve maximum accuracy, generalization and speed, we used only words that appeared in a polarity dictionary (from He, Lin and Alani, 2011) that appeared at least 50 times in the corpus and employed a pure discriminative classifier. We initially, tried libSVM from the Shogun machine learning package (Sonnenburg et. al. 2010). With it, we obtained accuracy of 89.7% on positive examples and 65.2% on negative examples, for an average of 77.1%. Given that performance on the negative examples was low, we switched to the Generalized Minimum Norm Problem SVM (Franc, 2005). This classifier produced accuracies of 87.6% on positive examples, 83.8% on negative examples, for an average performance of 85.7%, which was acceptable.

It is worth noting that some authors have reported accuracy in excess of 95% on this dataset using a combination generative graphical model plus support vector machine (SVM) approach (He, Lin & Alani, 2011). Part of the work involved in the contract under which this work was carried out was an attempt to recreate the high level of accuracy reported by He et al. Several attempts using different parameterizations failed to reproduce the very high accuracy reported by He et al, and as a result, those enhancements to the classifier were not implemented in the final product discussed here.

1.1.3 Evaluation

As an initial test, we tried the sentiment analysis mechanism on a set of reviews of the Amazon Kindle that were extracted from the Amazon website. It quickly became apparent that the classifier did not generalize well to these documents, with many misclassifications immediately obvious. We do not have ground truth labels with which to quantify the performance on this dataset, but it is clear that generalization across domains will be difficult. Note that we used a restricted vocabulary of polarity laden terms, which should help to improve generalization, and yet performance was disappointing. It may be necessary in the future to train classifiers for specific domains and/or to work to improve the vocabulary of polarity words.

The screenshot displays a web application interface for document management and sentiment analysis. At the top, there is a navigation bar with icons for Documents, Space, Time, and Help, along with a search bar containing "Sentiment102 AND 'Star Wars'", a "Go" button, and links for Reports and Manager. A user greeting "Welcome Ben Logout" is visible in the top right corner.

The main content area is divided into three sections:

- HANDLES**: A list of 1214 handles, with the first 14 shown. Each handle is color-coded and includes a count in parentheses. The handles are: star wars (143), special effects (75), phantom menace (63), obi wan (38), qui gon (29), george lucas (28), science fiction (26), star trek (25), starship troopers (20), obi wan kenobi (20), darth vader (19), liam neeson (17), toy story (16), anakin skywalker (16), star wars episode (14), natalie portman (13), star wars trilogy (13), and jake lloyd (13).
- Document Content**: Three document snippets are displayed, each with a title and a Doc ID. The first snippet is titled "cv953_6836p1" and has Doc ID: 1332421318.08.183. The second snippet is titled "cv953_6836p2" and has Doc ID: 1332421318.08.183. The third snippet is titled "cv953_6836p3" and has Doc ID: 1332421318.08.183. Each snippet contains text with sentiment analysis results, such as "star wars", "science fiction", "star wars trilogy", "PositiveSentiment", "special effects", "darth maul", "independence day", "star wars episode", "star trek", "harrison ford", "star wars trilogy", "george lucas", "audience to enjoy", "much has been and is still being made of the fact that mr.", "lucas has gone back and revised his trilogy for rerelease now", and "some fans".
- SEARCH OPTIONS**: A sidebar on the right containing search filters. It includes fields for "No of Documents" (200), "Max Document Word Size" (50), "No of Handles" (30), "Start Date", "End Date", and "Sentiments" (None, All (2), Other...). A "Less..." link is at the bottom.

Figure 1 Documents View, 67 "Star Wars" movie reviews have been split into document size 50, resulting in 1214 handles documents.



Figure 2 The Space View has been organized by Negative and Positive Sentiment.

Next we considered a query on the sentiment database itself. We chose to examine sentiment towards the movie *Star Wars*. The initial query returned 67 documents. HanDles focused mainly on the actors and miscellaneous issues. Consequently, we chose to divide these documents into shorter segments of 50 words. The resulting set of 1214 documents and 30 handles were much easier to interpret (see Figure 1). It is clear that dividing the documents is a critical capability. In the future, it would be worthwhile to investigate more sophisticated methods for division including segmenting at sentence boundaries.

Figure 2 shows the space view organized by sentiment. This was accomplished by “freezing” or locking the Positive Sentiment and Negative Sentiment handles in opposite regions in the display. These sentiment handles were then clicked repeatedly, or “pumped”, to pull related documents closer to them. As can be seen in Figures 3 through 5, sentiment can now be explored by clicking on the sentiment handles and other handles of interest in the display. We find that both “George Lucas” and “Obi Wan Kenobi” documents contain a lot of Negative Sentiment handles. Alternatively, “Special Effects” documents generally contained Positive Sentiment handles.

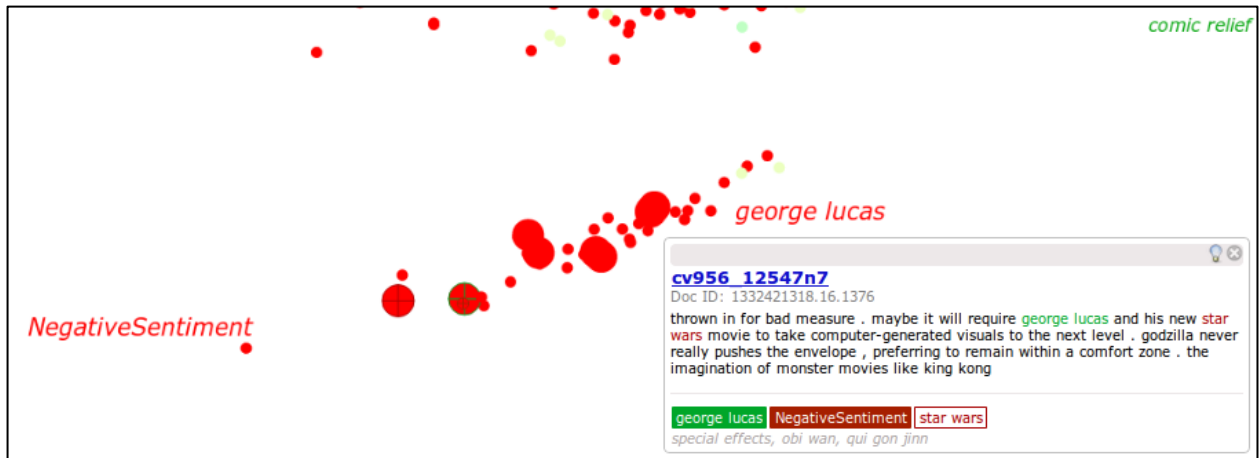


Figure 3 Negative Sentiment and “George Lucas”.

One of the difficulties that presents itself in when attempting to extract sentiment out movie review documents, is the context of the sentiment. A good example of this problem is displayed below in Figure 6. The text fragment that has been open reads:

“senate to make an appeal for justice . on tatooine , qui gon discovers a young slave boy , anakin skywalker (jake lloyd) , who not only can help them get the parts they need , but displays uncanny intelligence , insight , and instincts . qui gon senses the”

In this example the classifier has labeled the document as containing Positive Sentiment, and to some extent this is true because the descriptive terms “uncanny intelligence”, “insight” and “instincts” are both desirable and positive qualities to attribute to someone. The problem with classification lies in the context or role of the author. In this case, the author adopts the role or voice of “story teller”, rather than the voice of a movie reviewer whose sentiment and opinions we were initially interested in polling. While the problem of distinguishing the author's voice is more likely to occur in reviews of movie or books, it is still possible it could occur in other datasets.

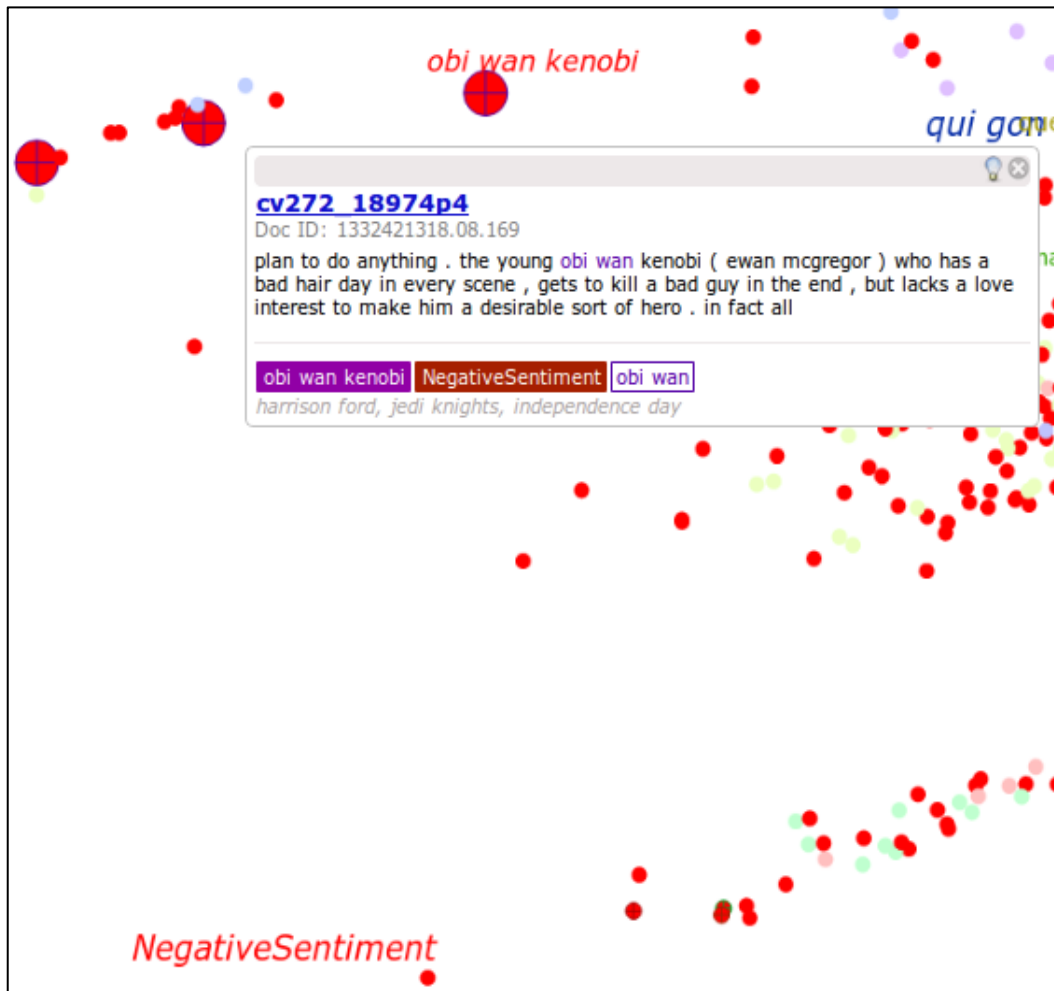


Figure 4 Negative Sentiment and “Obi Wan Kenobi”

Another potential problem that is unique to the handles interface is displayed in Figure 7. Ideally, a request to HanDles for “Star Wars” movie reviews would only return those reviews that are about the six Star Wars movies; however, this is not the case. There are many different movies reviewed in Pang and Lee's (2004) dataset, some of which are not reviewing a Star Wars movie, but do reference the Star Wars movies. For example, one of the reviews of the film Toy Story mentions a parody of Star Wars made by that movie. Another review compares the film, Starship Troopers to Star Wars saying that “Starship Troopers is very reminiscent [sic] of star wars , another kick-ass space opera”. While returning these documents is not necessarily a bad thing (the user may wish to know all references about a subject), it is worth the user bearing in mind that these types of documents may be returned by HanDles.

The Pang and Lee's (2004) dataset does not have dates associated with the individual movie reviews, so it is not possible to view sentiment in the Timegraph view. This was unfortunate, as viewing handles (of which Sentiment is one) over time is a very informative feature of the HanDles application.

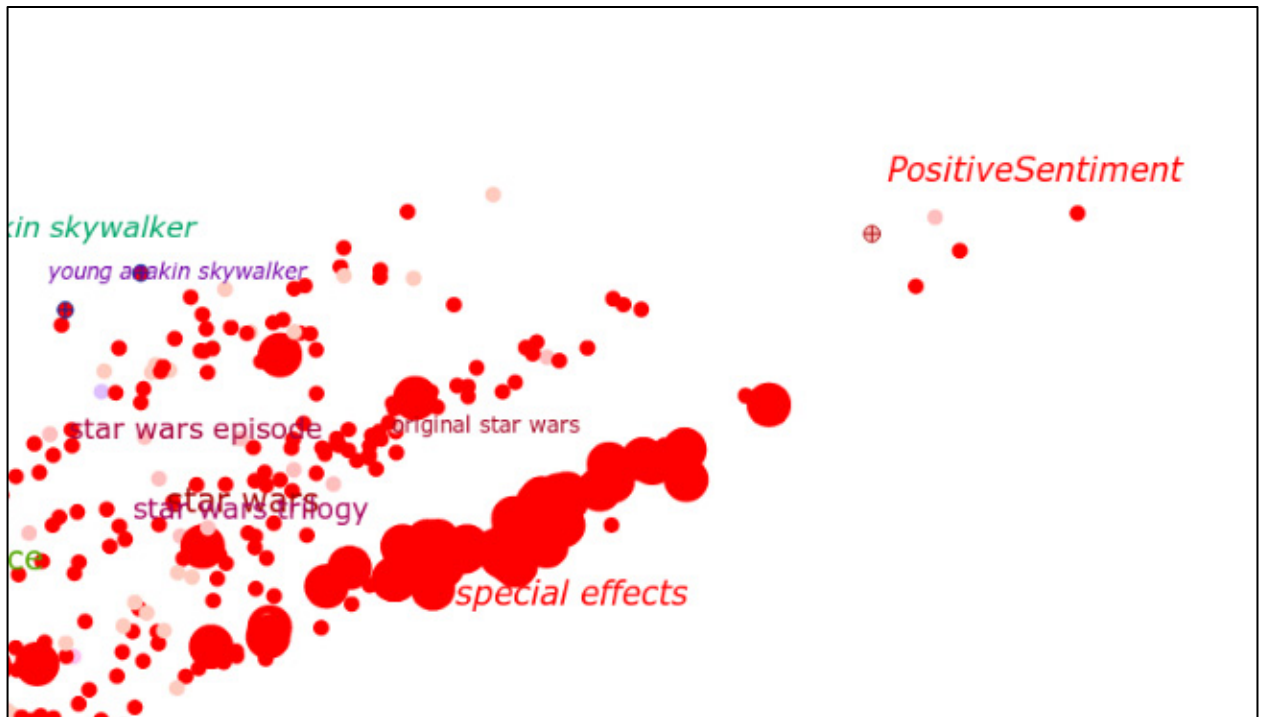


Figure 5 Positive Sentiment and “Special Effects”.

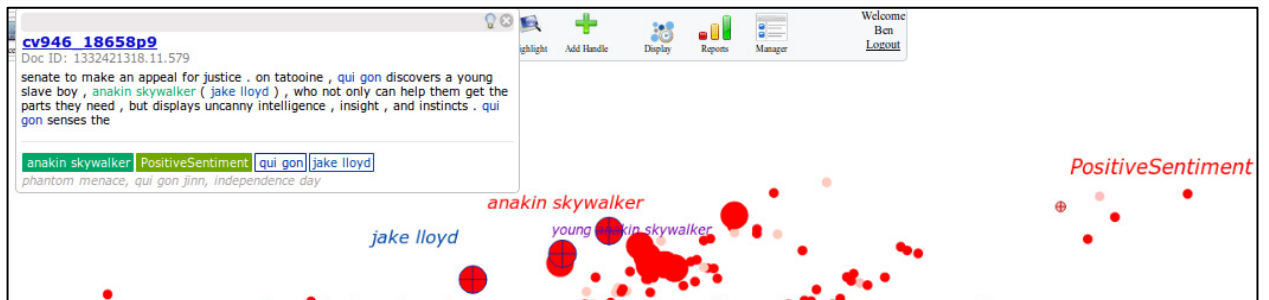


Figure 6 The potential problem of context for Sentiment Classification



Figure 7 The problem of “non Star Wars” reviews that mention “Star Wars”.

It should not come as a surprise that, in the above test of the capability, we obtained acceptable results when we used movie reviews taken from the training set. In a second test of the system, we selected 30 movie reviews written by IMDB users. Fifteen of the reviews were taken from a list of the worst movies ever made (*LOL*, *Jack And Jill*, *Meet The Spartans*, *Titanic*, *Dragonball*, *Epic Movie*, *Vampires Suck*, *Spy Kids*, *Disaster Movie*, *Catwoman*, *Bucky Larson*, *The Room*, *Battlefield Earth*, *Superbabies*, *The Hottie And The Nottie*) and 15 from a list of the best movies ever made (*Shawshank Redemption*, *The Dark Knight*, *Lord of the Rings*, *Inception*, *Se7en*, *Spirited Away*, *A Separation*, *Toy Story 3*, *The Lives Of Others*, *The Untouchables*, *Wall-E*, *Braveheart*, *Pan's Labrynth*, *Batman Begins*, *Gran Torino*). Figure 8 shows the film titles in the HanDles map. As a next step, we moves the PositiveSentiment and NegativeSentiment handles to opposite corners of the screen. Doing so causes the movie title on the screen handles to follow the most strongly associated sentiment handle. Then, to further separate the them, we continuously click or, “pump” the positive and negative sentiment handles to draw associated titles closer. As is clear in Figure 9, HanDles does an impressive job of differentiating good and poor reviews. HanDles did make two interesting misclassifications, however. It mistakenly classified the review for the Clint Eastwood film, *Gran Torino* as negative and the Adam Sandler flop, *Jack and Jill*, as positive (see Figure 10). An examination of the narrative in the reviews clarifies how the misclassification occurred. *Gran Torino* is a film about a grumpy, foul mouthed man and as such, the film’s description contains many terms that are generally associated with negative sentiment like, for example, foul, racist, dirty (as in *Dirty Harry*) racist, complaints, and garbage. In essence, the negative terms contained in the text misled the classifier.

“With his performance Eastwood shows you why people like himself, Jack Nicholson, or Paul Newman only come around once in a lifetime. Though Eastwood would rather focus on directing, he can still carry a movie with his on screen presence, and he's pure dynamite in "Gran Torino". Perhaps the poor box office results of "hollywoody" movies like Absolute Power, True Crime, Space Cowboys, and Blood Work, caused Eastwood to shy away from acting, but given cutting edge material to work with as "Million Dollar Baby" and "Gran Torino", he's as good as ever. His character as the racist and salty war vet makes you think of that old guy we've all had on our blocks with the garbage (sic) door open, the million tools hanging everywhere, and always fixing or building something. I found myself not wanting the movie to end because the scenes between himself and the various Hmong characters were priceless. There may be complaints over the racist remarks and scenes, but

Eastwood pulls it off in a way a real person like that would talk or act to a point where it ends up being lighthearted. I'm not going to give the plot away, but if you like your Clint Eastwood as a hard-nosed tough guy with foul language alla Dirty Harry or Heartbreak Ridge, you'll love this movie!!"

On the other hand, the film, *Jack and Jill* was classified as containing positive sentiment, despite the review being clearly scathing.

"I will start by saying that I have enjoyed many Adam Sandler movies and have found him to be a generally funny guy when I've seen him interviewed and when he was on Saturday Night Live. I laughed gleefully through Anger Management, Mr. Deeds, Billy Madison, 50 First Dates, and Happy Gilmore. Funny People and Just Go With It were awesome movies! I was brought to tears in Sandler's emotional portrayal in Reign Over Me. I have great respect for the man as a comedian and actor. But Jack and Jill is abysmal. The "jokes" are not only bizarrely misplaced - THEY ARE NOT FUNNY. I did not hear a single laugh, not even a slight giggle from any audience member in the theater. In fact, almost 1/3 of them walked out before it was over. Those who stayed, openly derided the flick as we all exited the theater in utter disgust and sadness. I don't know why Al Pacino was in this movie, his acting made it seem like he was forced at gun point to do this movie. Nick Swardson and Tim Meadows are way too funny to be in such a disaster. Especially given Swardson's stellar performance in Just Go With It. This movie is not a flop, its not an "oops", its not a mistake - it's a career ending pile of trash. A career ender that started with Sandler's god awful "Grown Ups" and climaxes with this revolting hunk of garbage. Sorry Adam, it's over."

Here again, the film's review is clearly negative, but the document contains several positive terms like, funny, laughed, gleefully, Happy, respect, funny (interestingly, "funny" appears in the phrase "not funny"), laugh, giggle, and stellar. A document's classification is made by weighing its positive and negative aspects, and in this case, there was clearly enough positive sentiment expressed in the review to mislead the classifier.



Figure 8 HandDles map of 30 film reviews from IMDB users

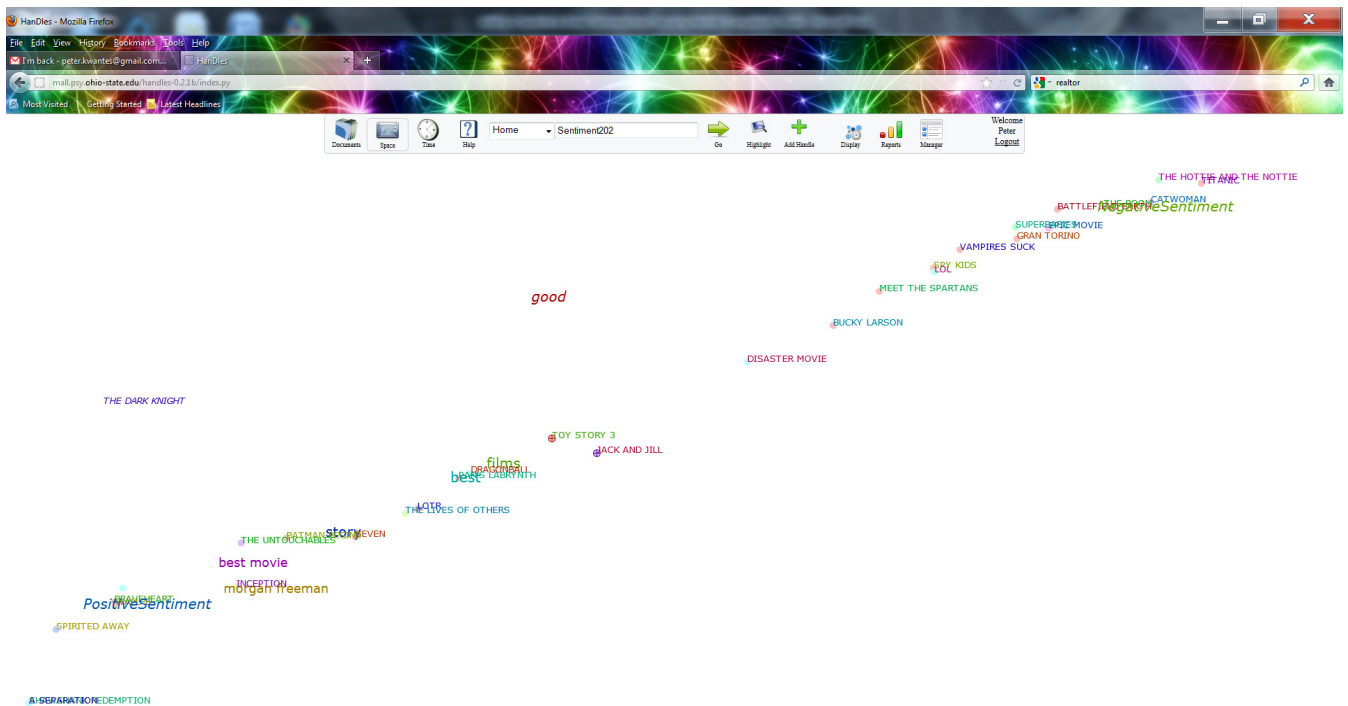


Figure 9 Alignment of the film titles when positive and negative handles are pulled apart and 'pumped'

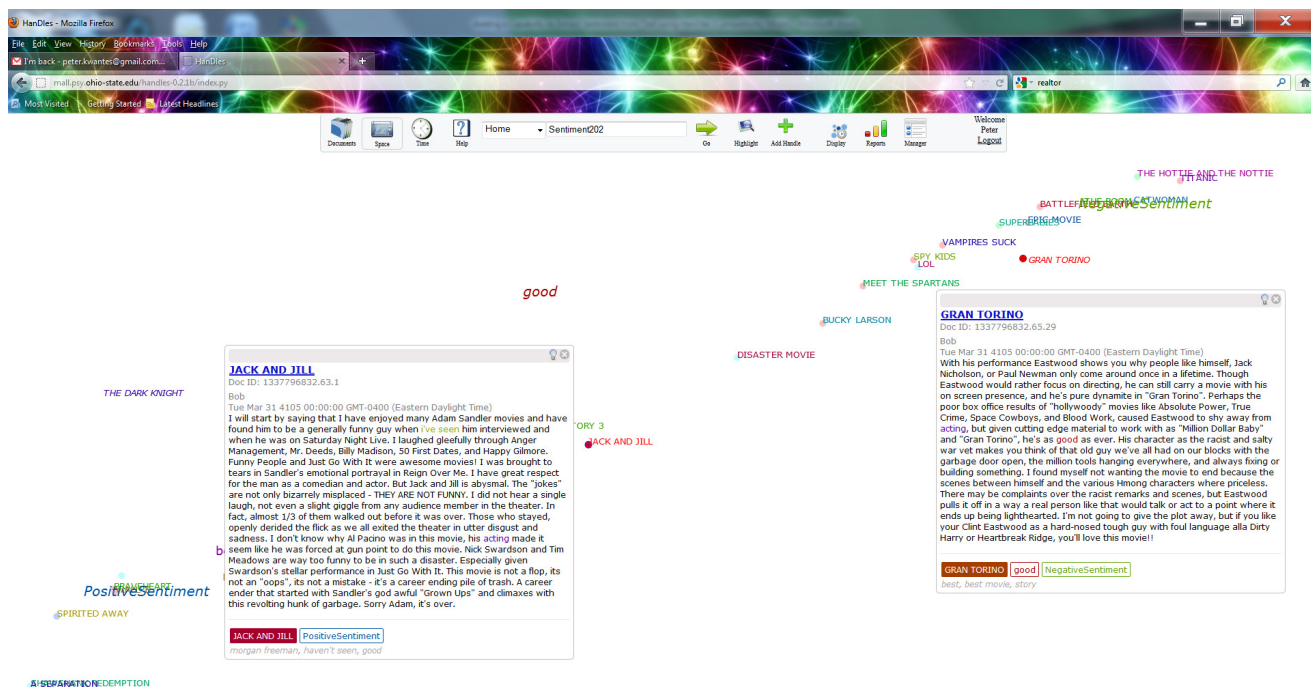


Figure 10 HanDles misclassified two reviews: Gran Torino and Jack and Jill.

2 Conclusion and Recommendations

The most important finding presented in this report is that it will be necessary in the future to train classifiers for specific domains and/or to work to improve the vocabulary of polarity words used by HanDles or any other mechanism that could be used to automatically classify documents according to their sentiment. In particular, the classifier should be trained on documents that are similar to those which will be examined in operational settings. For example, in an Influence Operations context, CF analysts may be interested in measuring positive and negative sentiment expressed by locals toward local governance or the Canadian Forces. Movie reviews do not provide the ideal class of document with which to train a sentiment analysis system. Instead, a large document collection comprising text written for and by members of the local population would be most appropriate. To this end, it is important to decide at the beginning of an influence operation, what kinds of documents from the area of operations should be gathered and submitted to the system.

Another conclusion from the work reported here is that the ability to divide the documents is a critical capability of Sentiment Analysis in HanDles, so future development work in this area would be of benefit to the HanDles system. Traditional sentiment analysis applies positive and negative classifications to whole documents. As we demonstrated above however, such a coarse level of analysis will often fail to capture what aspects of the topics under discussion are written about positively or negatively. For example, the review of *Jack and Jill* above is decidedly negative, but contains positive sentiment when discussing its cast. HanDles' ability to subdivide documents allows it to detect the sentiment associated with various topics discussed in a document—a capability not common to tools of this sort, but nonetheless critical for accurate situational awareness and measurement of a campaign's effectiveness in operational settings.

HanDles is ready for exploitation. The next step in its transition to operational use will be the generation of an appropriate collection of documents to use for training the system to distinguish documents expressing positive and negative sentiment. To do so however, DRDC and CF personnel from the Intelligence and Influence Activities context will need to work together to decide on where those documents will come from.

Sentiment is one of several kinds of classification that can be applied to documents. In future, we can envisage other forms of discrimination being introduced to the tool. For example, the SVM in HanDles could be trained to automatically distinguish threatening from non-threatening narrative in blogs or violent from non-violent events described in intelligence or situation reports. The point is that the SVM is a generic classifier, and can be trained to classify documents on any dimension so long as the documents in the training set have ground-truth labels.

References

- [1] Bo Pang and Lillian Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In Proceedings of the ACL, 2004.
- [2] Yulan He, Chenghua Lin, and Harith Alani. Automatically extracting polarity-bearing topics for cross-domain sentiment classification. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1, HLT '11, pages 123–131, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics
- [3] Vojtech Franc. Optimization Algorithms for Kernel Methods. Research Report. CTU-CMP-2005-22. CTU FEL Prague. 2005. <ftp://cmp.felk.cvut.cz/pub/cmp/articles/franc/Franc-PhD.pdf>. 2005.
- [4] Soeren Sonnenburg, Gunnar Raetsch, Sebastian Henschel, Christian Widmer, Jonas Behr, Alexander Zien, Fabio de Bona, Alexander Binder, Christian Gehl, and Vojtech Franc. The SHOGUN Machine Learning Toolbox. Journal of Machine Learning Research, 11:1799-1802, June 2010.

This page intentionally left blank.

List of symbols/abbreviations/acronyms/initialisms

DND	Department of National Defence
DRDC	Defence Research & Development Canada
DRDKIM	Director Research and Development Knowledge and Information Management
R&D	Research & Development
OMSA	Opinion Mining and Sentiment Analysis
SVM	Support Vector Machine
IMDB	Internet Movie Database

This page intentionally left blank.

DOCUMENT CONTROL DATA		
(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)		
1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)	2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)	
Defence R&D Canada – Toronto 1133 Sheppard Avenue West P.O. Box 2000 Toronto, Ontario M3M 3B9	UNCLASSIFIED (NON-CONTROLLED GOODS) DMC A REVIEW: JUNE 2010	
3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)		
Adding a Capability to Extract Sentiment from Text using HanDles:		
4. AUTHORS (last name, followed by initials – ranks, titles, etc. not to be used)		
Simon Dennis; Benjamin Stone; Jihun Hamm; Peter Kwantes		
5. DATE OF PUBLICATION (Month and year of publication of document.)	6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)	6b. NO. OF REFS (Total cited in document.)
May 2012	28	5
7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)		
Technical Memorandum		
8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)		
Defence R&D Canada – Toronto 1133 Sheppard Avenue West P.O. Box 2000 Toronto, Ontario M3M 3B9		
9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)	
15ah02	W7711-088147/001/TOR	
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)	
DRDC Toronto TM 2012-063		
11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)		
Unlimited		
12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.)		
Unlimited		

13. **ABSTRACT** (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

HanDles is a document visualization tool developed by Ohio State University for DRDC Toronto. One aspect of documents that might be of interest to analysts is the extent to which they express positive or negative opinion or sentiment toward some issue or group. In this report, we describe how HanDles was extended to include the ability to classify documents as containing predominantly positive or negative sentiment. To do so, we trained the semantic model underlying HanDles' understanding of the document collection to distinguish positive from negative documents. Our tests of the system suggested that its ability to discriminate positive from negative documents would be greatly improved by selecting a training collection that is similar in nature and content to the documents that will be evaluated in operational settings.

HanDles est un outil **de visualisation** de documents conçu par l'Ohio State University pour RDDC Toronto. Une caractéristique des documents qui peut s'avérer intéressante pour les analystes est l'importance de l'opinion positive ou négative que dégagent ces documents à l'égard de certaines questions ou de certains groupes. Dans le présent rapport, nous décrivons comment nous avons amélioré HanDles afin qu'il prenne en charge la classification de documents selon la prédominance, dans leur contenu, de sentiments positifs ou négatifs. Pour ce faire, nous avons formé la compréhension du modèle sémantique sous-jacent à HanDles quant au recueil de documents utilisé afin qu'il soit en mesure de distinguer les documents positifs de ceux qui sont négatifs. Nos essais du système nous poussent à croire qu'il est possible d'accroître considérablement sa capacité à différencier les documents selon le sentiment qui s'en dégage en choisissant un recueil de formation dont la nature et le contenu ressemblent à ceux des documents qui seront évalués dans un contexte opérationnel.

14. **KEYWORDS, DESCRIPTORS or IDENTIFIERS** (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Handles; document visualization; sentiment analysis; opinion mining

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale



www.drdc-rddc.gc.ca

